

Pheet meets C++11

Manuel Pöter *

November 10, 2014

Abstract

Pheet is a C++ task-scheduling framework that allows for easy customization of internal data-structures. The implementation was started before the C++11 standard was committed and therefore did not use the new standardized memory model but compiler/platform specific intrinsics for atomic memory operations. This not only makes the implementation harder to port to other compilers or architectures but also suffers from the fact that prior C++ versions did not specify any memory model.

In this report I discuss the porting of one of the internal Pheet data structures to the new memory model and provide reasoning about the correctness based on the semantics of the memory consistency model. Using two benchmarks from the Pheet benchmark suite I compare the performance of the original against the new implementation which shows a significant speedup under certain conditions on one of the two test machines.

Introduction

Pheet¹, developed by Wimmer [10], is a highly customizable task scheduling framework developed in C++ that allows comparison between different implementations of data structures used in the scheduler, as well as comparisons between entirely different schedulers. The customization options are based on the powerful template meta-programming possibilities provided by the C++ language. Even though the framework is fully customizable and allows replacing almost all of the internal data structures, the template meta-programming approach produces high-performance code since C++ templates are instantiated at compile time and therefore all the usual compiler optimizations like inline expansion, dead-code elimination, etc. can be applied. It is easy to use and provides good performance.

Although the framework makes extensive use of many new language features that were introduced in C++11 (mainly template aliases and variadic templates) the development started before the C++11 standard was officially adopted by the standards committee.

*This report was the outcome of a project from autumn 2013 to summer 2014 at the Parallel Computing group at the Vienna University of Technology. The project was supervised by Martin Wimmer and Jesper Larsson Träff, and also resulted in [12].

¹<http://www.pheet.org>

For this reason the original implementation did not use the new memory model that was introduced with C++11, but instead relied on macros that resolved to compiler/platform specific intrinsics.

The C++ Standard prior to C++11 specified program execution in terms of observable behavior, which in turn described sequential execution on an implicitly single-threaded abstract machine. Therefore multi-threaded C++ programs relied on set of libraries for threading support like POSIX threads, Win32, or Boost. Unfortunately a pure library approach, in which the compiler is designed independently of threading issues, includes all sorts of problems [4]. Without a clearly defined memory model as a common ground between the compiler, the hardware, the threading library, and the programmer, multi-threaded C++ code is fundamentally at odds with compiler and processor-level optimizations [8].

In August 2011 the ISO C++ committee finally approved the next C++ standard – commonly referred to as C++11. This new standard defines a multi-threaded abstract machine, together with a well-defined memory model and library support for interaction between threads. The memory model defines when multiple threads may access the same memory location, and specifies when updates by one thread become visible to other threads. This not only allows development of multi-threaded applications in a platform independent and portable way, but also to formally prove the correctness based on the memory order semantics.

Unfortunately at the beginning of the development of Pheet the C++11 standard was not finalized and no implementation of the proposed memory model was available in gcc. For this reason none of the early data structures were developed using the new memory model, but instead with a few macros that resolve to compiler specific intrinsics for atomic operations like fetch-and-add or compare-and-swap. My task was to provide an adapted implementation of the CentralKStrategy data structures based on the new memory model instead of the old macros and to provide a reasoning about the correctness of the new implementation based on the memory order semantics.

The following description of the implementation and the code listings are based on revision 573 from <https://launchpad.net/pheet>.

The C++11 memory model²

The new memory model of C++11 is largely based on the work by Boehm, Alexandrescu et al [5, 3].

One of the most important aspects is the definition of a data race [6, p. 14]:

The execution of a program contains a data race if it contains two conflicting actions in different threads, at least one of which is not atomic, and neither happens before the other. Any such data race results in undefined behavior.

Where conflicting actions are defined as follows [6, p. 11]:

²Since the official C++ standard is not freely available I will instead refer to the "Working Draft, Standard for Programming Language C++" from January 2012 [6] which contains the C++11 standard plus minor editorial changes.

Two expression evaluations conflict if one of them modifies a memory location and the other one accesses or modifies the same memory location.

This definition implies that any program written according to the old standard that uses some other threading libraries and shares any data between those threads exhibits undefined behavior. The memory operations are ordered by means of the *happens-before* relationship that can be roughly described as follows:

Let A and B represent operations performed by a multi-threaded process. If A happens-before B, then the memory effects of A effectively become visible to the thread performing B before B is performed.

The *happens-before* relation (denote: \rightarrow) is a strict partial order and as such it is transitive, irreflexive and antisymmetric.

transitivity - $\forall a, b, c$, if $a \rightarrow b$ and $b \rightarrow c$, then $a \rightarrow c$

irreflexivity - $\forall a, a \not\rightarrow a$

antisymmetry - $\forall a, b$, if $a \rightarrow b$ then $b \not\rightarrow a$

The complete formal definition specifically for C++ can be found in [6, p. 11-13].

A *happens-before* order between two operations from the same thread (program order) is implicitly given by the *sequenced-before* order [6, p. 10,13]. A *happens-before* order between two operations from different threads (in the standard this is referred to as *inter-thread-happens-before*) must be established using atomic operations.

The following memory orders can be specified for each atomic operation (from strong to relaxed):

- `memory_order_seq_cst`
- `memory_order_acq_rel`
- `memory_order_release`
- `memory_order_acquire`
- `memory_order_consume`
- `memory_order_relaxed`

`memory_order_consume` and `memory_order_acquire` can only be used for operations that perform a *read*, `memory_order_release` can only be used for operations that perform a *write* and `memory_order_acq_rel` can only be used for operations that perform a *read-modify-write* operation. Although the language does not enforce these constraints some implementations do check them at runtime³.

A *happens-before* relationship can be established by using the following combinations of memory orders⁴:

³For example the Microsoft STL implementation verifies these constraints at runtime when the `DEBUG` macro is defined.

⁴`memory_order_acq_rel` is the combination of `memory_order_release` and `memory_order_acquire`. So wherever either one is used it is also possible to use `memory_order_acq_rel`.

- `memory_order_seq_cst + memory_order_seq_cst`
- `memory_order_acquire + memory_order_release`
- `memory_order_consume + memory_order_release`

There is always a *happens-before* order between two `memory_order_seq_cst` operations [6, p. 1100].

An atomic operation A that performs a store-release operation on an atomic object M *synchronizes with* an atomic operation B that performs a load-acquire operation on M and takes its value from any side effect in the release sequence (defined below) headed by A . This *synchronize-with* order is compatible with the *inter-thread-happens-before* order. The *dependency-ordered-before* relation resulting from `memory_order_consume` operations is more complicated, but since it is not used in the ported implementation I will not go into more details (these can be found in [6, p. 12]). `memory_order_relaxed` can never be used to create a *happens-before* order.

All modifications to a particular atomic object occur in some particular total order, called the *modification order*. If A and B are modifications of an atomic object M and A *happens-before* B , then A precedes B in the modification order of M . There are separate modification orders for each atomic object and there is no requirement that these can be combined into a single total order for all objects.

A *release sequence* is a subsequence of the modification order of an atomic object. It is headed by a release operation A and followed by an arbitrary number of

- atomic operations performed by the same thread that performed A or
- atomic read-modify-write operations.

For operations performed by the same thread that performed A it is not relevant which memory order is used – even `memory_order_relaxed`.

The standard describes two different compare-and-swap operations for atomic objects - `compare_exchange_strong` and `compare_exchange_weak`. The difference is that `compare_exchange_weak` is allowed to fail spuriously, that is, act as if `*obj != *expected` even if they are equal, but can result in better performance on some platforms.

CentralKStrategy

CentralKStrategy is a task storage data structure that is internally used by the Pheet scheduler. It creates a global priority ordering between all the tasks available in the system, while allowing each worker thread to miss up to k of the newest tasks, as long as each task is seen by at least one thread. This k -relaxation, which is based on the scheme introduced by Afek et al. [1], is used to improve scalability of the data structure because it is shared by all threads.

The implementation is split into two components. The first is a global shared array that is used to share tasks between all threads and to maintain information about which tasks have to be globally visible, to avoid violations of the k -relaxation guarantees. The second is a local priority queue for each thread (in Pheet also referred to as "place") that is used to maintain the priority ordering of the task visible to each thread.

A detailed description of the CentralKStrategy data structure can be found in [13] and [10, pp. 81].

The CentralKStrategy implementation consists of the following data structures:

- CentralKStrategyTaskStorageItem
- CentralKStrategyTaskStorageImpl
- CentralKStrategyTaskStoragePlace
- CentralKStrategyTaskStorageDataBlock

My first task was to identify all the members of these data structures that have to be defined as atomics in order to avoid any data races. Once these members were changed to atomics the implementation would be correct since all the operations on atomics use `memory_order_seq_cst`. However, from a performance point of view this is far from optimal, so for every single operation on any atomic member I relaxed the used memory order as far as possible, while ensuring that all the required *happens-before* relations are kept in place.

CentralKStrategyTaskStorageItem

This data structure represents an item that is referenced from the global array as well as from the thread's local priority queue. Each item represents a generated task to be executed by some worker thread. It has the following members:

`Pheet::Scheduler::BaseStrategy* strategy` - a pointer to a *scheduling strategy* that is used internally by the scheduler. More details about how the scheduler can be configured with strategies can be found in [11].

`TT data` - some implementation specific data.

`size_t position` - the position of this item in the global array, also used to mark this item as "taken".

`size_t orig_position` - the position in the global array where this item was originally inserted.

`Place* owner` - pointer to the `CentralKStrategyTaskStoragePlace` instance that owns this item.

`void (Place::*item_push)(Self* item, size_t position)` - pointer to a member function that is used to push the item to a thread's local priority queue.

It does not contain any relevant methods.

The only member that can participate in a data race is `position`. All other members are set by the owning thread before the item is made globally available (by storing a reference in the global array) and the other threads only read these members. Therefore it is sufficient to make the member `position` atomic.

position

position is set to the same value as **orig_position** during initialization (together with all the other members) before the item is made globally visible. All other accesses are either loads followed by a **compare_exchange** operation to update **position** and mark the item as "taken" (**position** != **orig_position**), or loads followed by a comparison to check whether the item is still active.

In either case a load that "sees" an old value is not an issue – either the following **compare_exchange** operation fails or the item is still seen as active which may result in some additional work (e.g. the item is inserted in the thread local heap structure), but does not affect correctness. Since there are no interdependencies with other values there is no need for a *synchronize-with* relation between any loads and stores, so all operations can use **memory_order_relaxed**.

CentralKStrategyTaskStorageImpl

This data structure represents the global shared array. There exists a single instance that is shared by all places (threads). It has the following members:

size_t tail - the global tail index in the global array – **head** indexes are thread local. It is guaranteed that for every thread all the items between the threads local **head** and the global **tail** are filled with items.

DataBlock* start_block - pointer to the first data block of the global array (an instance of **CentralKStrategyTaskStorageDataBlock**).

size_t num_places - the total number of places (threads) that are using this data structure (must be known in advance).

It does not contain any relevant methods, most of the methods are implemented in **CentralKStrategyTaskStoragePlace**.

tail can be accessed concurrently by several threads and therefore has to be defined as atomic. **start_block** and **num_places** are only set during initialization and therefore cannot participate in any data race.

tail

While every thread has its own local **head** variable (see description of the data structure **CentralKStrategyTaskStoragePlace**) there is only a single global tail. It is guaranteed that all the entries between some **head** and the global **tail** contain valid pointers to items. Entries are inserted behind **tail** according to the *k* relaxation. As soon as all *k* entries behind **tail** are used up, **tail** is updated accordingly.

In order to guarantee a *happens-before* order for load/store accesses to the entries the **compare_exchange_weak** operation on **tail** in **push** must use **memory_order_release** and the load operation in **update_heap** must use **memory_order_acquire**. Both methods are implemented in **CentralKStrategyTaskStoragePlace**.

It is sufficient to use **compare_exchange_weak** for the update since it is executed in a loop until either this or some other thread is successful, so a spurious failure does not cause any problem (see Listing 1).

Listing 1: Update of `tail` in `CentralKStrategyTaskStoragePlace::push`

```
1 size_t nold_tail = task_storage->tail.load(std::memory_order_relaxed);
2 ptrdiff_t diff = static_cast<ptrdiff_t>(cur_tail) - static_cast<ptrdiff_t>(
    nold_tail);
3 while (diff > 0)
4 {
5     if (task_storage->tail.compare_exchange_weak(nold_tail, cur_tail, std::
        memory_order_release, std::memory_order_relaxed))
6         break;
7     diff = static_cast<ptrdiff_t>(cur_tail) - static_cast<ptrdiff_t>(
        nold_tail);
8 }
```

CentralKStrategyTaskStoragePlace

This data structure contains all the thread local data – including the thread local priority queue. It has the following members:

`PerformanceCounters pc` - a class containing some data structure specific performance counters.

`TaskStorage* task_storage` - a pointer to the `CentralKStrategyTaskStorageImpl` singleton instance that is shared by all threads.

`StrategyRetriever sr` - the strategy retriever that is used in the heap (local priority queue).

`StrategyHeap heap` - the thread's local priority queue.

`DataBlock* tail_block` - pointer to the first data block in the linked list (that builds the global shared array) that this thread knows of.

`DataBlock* head_block` - pointer to the last data block in the linked list (that builds the global shared array) that this thread knows of.

`size_t head` - the thread's local head index in the global shared array.

`ItemMemoryManager items` - memory manager for the items created by this thread.

`DataBlockMemoryManager data_blocks` - memory manager for the data blocks created by this thread.

Since every thread has its own `CentralKStrategyTaskStoragePlace` instance and does not touch instances from other threads the members of this data structure cannot participate in any data race and are not required to be atomic.

CentralKStrategyTaskStorageDataBlock

The global array is actually implemented as a linked list of arrays. This data structure represents one such part of the global array. It has the following members:

`Item* data[BlockSize]` - an array of `CentralKStrategyTaskStorageItem` pointers. `BlockSize` is a template parameter and can therefore be tuned.

`size_t offset` - the index offset for this block (the first element in this block has index `offset` in the global array).

`DataBlock* next` - pointer to the next data block in the linked list that builds the whole global array.

`size_t active_threads` - the number of threads that have not yet discarded this block.

`bool active` - a flag that signals whether this block is currently in use or not.

The only member that cannot participate in a data race is `offset` since it is only set during the initialization of a data block. All other members can be accessed concurrently and therefore have to be atomics.

Listing 2: C++11 code of `CentralKStrategyTaskStorageDataBlock::deregister`

```
1 void deregister()
2 {
3     size_t old = active_threads.fetch_sub(1, std::memory_order_relaxed);
4     if (old == 1)
5     {
6         // cleanup data items
7         // ...
8         next.store(nullptr, std::memory_order_relaxed);
9         active.store(false, std::memory_order_release);
10    }
11 }
```

A block can be reused as soon as all tasks (items) from it have been executed and the local `head` indexes of all threads point to positions in successor blocks. A kind of reference counting system is used to track whether a block can be reused or not. During initialization of the block `active` is set to `true` and `active_threads` is set to the number of total threads. When a thread increases its local `head` index to point to a new block it calls `deregister` on the previous block. `deregister` decreases `active_threads` to signal that this thread will no longer reference this block (see Listing 2). If `active_threads` drops to zero all resources allocated by this block are released and `active` is set to `false` to signal that this block can now be reused.

active

`active` is used to mark whether a block is currently in use (i.e. can be accessed by several threads concurrently) or not. It is accessed in the following methods:

`is_reusable` - checks if `active` is false (using `memory_order_relaxed`)

`add_block` - reads `active` (using `memory_order_acquire`) and sets `active` to `true` (using `memory_order_relaxed`)

`deregister` - sets `active` to `false` (using `memory_order_release`)

`is_reusable` (see Listing 3) and `add_block` are always executed by the thread that owns the data block. `deregister` can be executed by several threads concurrently, however only the last thread that decrements `active_threads` will clean up the data block and eventually set `active` to `false`.

Listing 3: `CentralKStrategyTaskStorageDataBlock::is_reusable`

```
1 bool is_reusable() const
2 {
3     return !active.load(std::memory_order_relaxed);
4 }
```

We need to ensure a *happens-before* order between resetting the items in `deregister` and subsequent accesses to them when the block is getting reused after `add_block` added it to the linked list. The load operation in `is_reusable` uses `memory_order_relaxed` to avoid potentially expensive acquire operations in cases where `active` is still `true` and the block cannot be reused yet. Instead `add_block` starts with a load operation using `memory_order_acquire` (see Listing 4). The load operation is solely required to establish the *synchronize-with* relation with the store in `deregister`. A *synchronize-with* relation is only established when the acquire load takes the value that has been written by the release store. In `add_block` the value returned by the load can safely be ignored since it is guaranteed to be `false` (the value written in `deregister`) because `add_block` for this block is only called when `is_reusable` returned `true` – which (as can be seen in Listing 3) is only the case when `active` is `false`.

Listing 4: `CentralKStrategyTaskStorageDataBlock::add_block`

```
1 void add_block(Self* block, size_t num_places)
2 {
3     // establish synchronize-with relationship with
4     // active.store from deregister()
5     block->active.load(std::memory_order_acquire);
6     block->active_threads.store(num_places, std::memory_order_relaxed);
7     block->active.store(true, std::memory_order_relaxed);
8
9     Self* pred = this;
10    block->offset = pred->offset + BlockSize;
11    auto nextBlock = pred->next.load(std::memory_order_relaxed);
12    while (nextBlock == nullptr)
13    {
14        if (pred->next.compare_exchange_weak(nextBlock, block,
15                                             std::memory_order_release,
16                                             std::memory_order_relaxed))
17            return;
18    }
19    // we failed to add the block, but some other thread must have succeeded
20    // -> make our block reusable again
21    block->active.store(false, std::memory_order_relaxed);
22 }
```

It would also be possible to create the *synchronize-with* relation between `deregister` and `is_reusable` by using an explicit `atomic_thread_fence` with `memory_order_acquire` (see Listing 5).

This way the *synchronize-with* relation with the release store in `deregister` would only be established when `active` is actually false. Therefore `deregister` *happens-before* `is_reusable` and `is_reusable` *happens-before* `add_block`. From the transitivity of the *happens-before* order follows that `deregister` *happens-before* `add_block`. This way the additional load operations could be omitted. Unfortunately I came up with this idea only a while after I finished the implementation.

Listing 5: Alternative implementation of `is_reusable` using an explicit `atomic_thread_fence`

```

1 bool is_reusable() const
2 {
3     bool result = !active.load(std::memory_order_relaxed);
4     if (result)
5         std::atomic_thread_fence(std::memory_order_acquire);
6     return result;
7 }

```

active_threads

`active_threads` is initialized with the number of total threads in `add_block` and updated using `fetch_sub` in `deregister`. Since this member is only used to track whether there are still any threads that may read from this block, but it is not used to order other accesses, `memory_order_relaxed` can be used for both operations.

next

The shared global array is implemented as a single-linked list of data blocks that are linked via the `next` pointer. The update of the `next` pointer via `compare_exchange_weak` in `add_block` must use `memory_order_release` and the according loads in the methods `push` and `deregister_old_blocks` of `CentralKStrategyTaskStoragePlace` must use `memory_order_acquire` to establish a *synchronize-with* relation.

The transitivity of *happens-before* guarantees that resetting the values in `deregister` *happens-before* any load of an entry from a reused data block in `push` - established via the *synchronize-with* relation between the load/store pairs of `active` and `next`.

It suffices to use `compare_exchange_weak` for the update since it is executed in a loop until either this thread or some other thread is successful, so a spurious failure does not cause any problem (see Listing 4).

data

For the `compare_exchange_strong` operation on `data` in the `put` method it is required to use `memory_order_release` to ensure a *happens-before* relation between initialization and subsequent loads of the members of `CentralKStrategyTaskStorageItem`.

In this case it is necessary to use `compare_exchange_strong` to avoid spurious failures because the thread tries to update each entry of a block with k entries only once (see Listing 6). If this attempt fails for all k entries then it is assumed that all entries are

already in use and `tail` is updated accordingly. However, in the face of spurious failures it may happen that the assumption that all entries are already set is no longer correct. This would break the guarantee that all entries between `tail` and some `head` contain valid references.

Listing 6: `CentralKStrategyTaskStorageDataBlock::put`

```

1 bool put(size_t& cur_tail, Item* item)
2 {
3     const size_t k = item->strategy->get_k();
4     size_t array_offset = cur_tail - offset;
5     while (array_offset < BlockSize)
6     {
7         const size_t cur_k = std::min(k, BlockSize - array_offset - 1);
8         const size_t to_add = Pheet::template rand_int<size_t>(cur_k);
9         const size_t i_limit = to_add + std::min(Tests, cur_k + 1);
10        for (size_t i = to_add; i != i_limit; ++i)
11        {
12            const size_t wrapped_i = i % (cur_k + 1);
13            auto& elem = data[array_offset + wrapped_i];
14            if (elem.load(std::memory_order_relaxed) == nullptr)
15            {
16                auto position = cur_tail + wrapped_i;
17                item->orig_position = position;
18                item->position.store(position, std::memory_order_relaxed);
19                Item* null = nullptr;
20                if (elem.compare_exchange_strong(null, item,
21                    std::memory_order_release,
22                    std::memory_order_relaxed))
23                    return true;
24            }
25        }
26        cur_tail += cur_k + 1;
27        array_offset = cur_tail - offset;
28    }
29    return false;
30 }

```

The diagram in Figure 1 shows the two classes `CentralKStrategyTaskStoragePlace` and `CentralKStrategyTaskStorageDataBlock` together with all their methods that access any atomic variables and calls to other relevant methods. To provide a more complete picture of what is happening the methods of `CentralKStrategyTaskStorageDataBlock` are "inlined". The different variables are highlighted in different colors so that all the places where they are used can be easily spotted. It also shows the calls between the different methods as well as the established *synchronize-with* relations.

Pheet CentralKStrategy

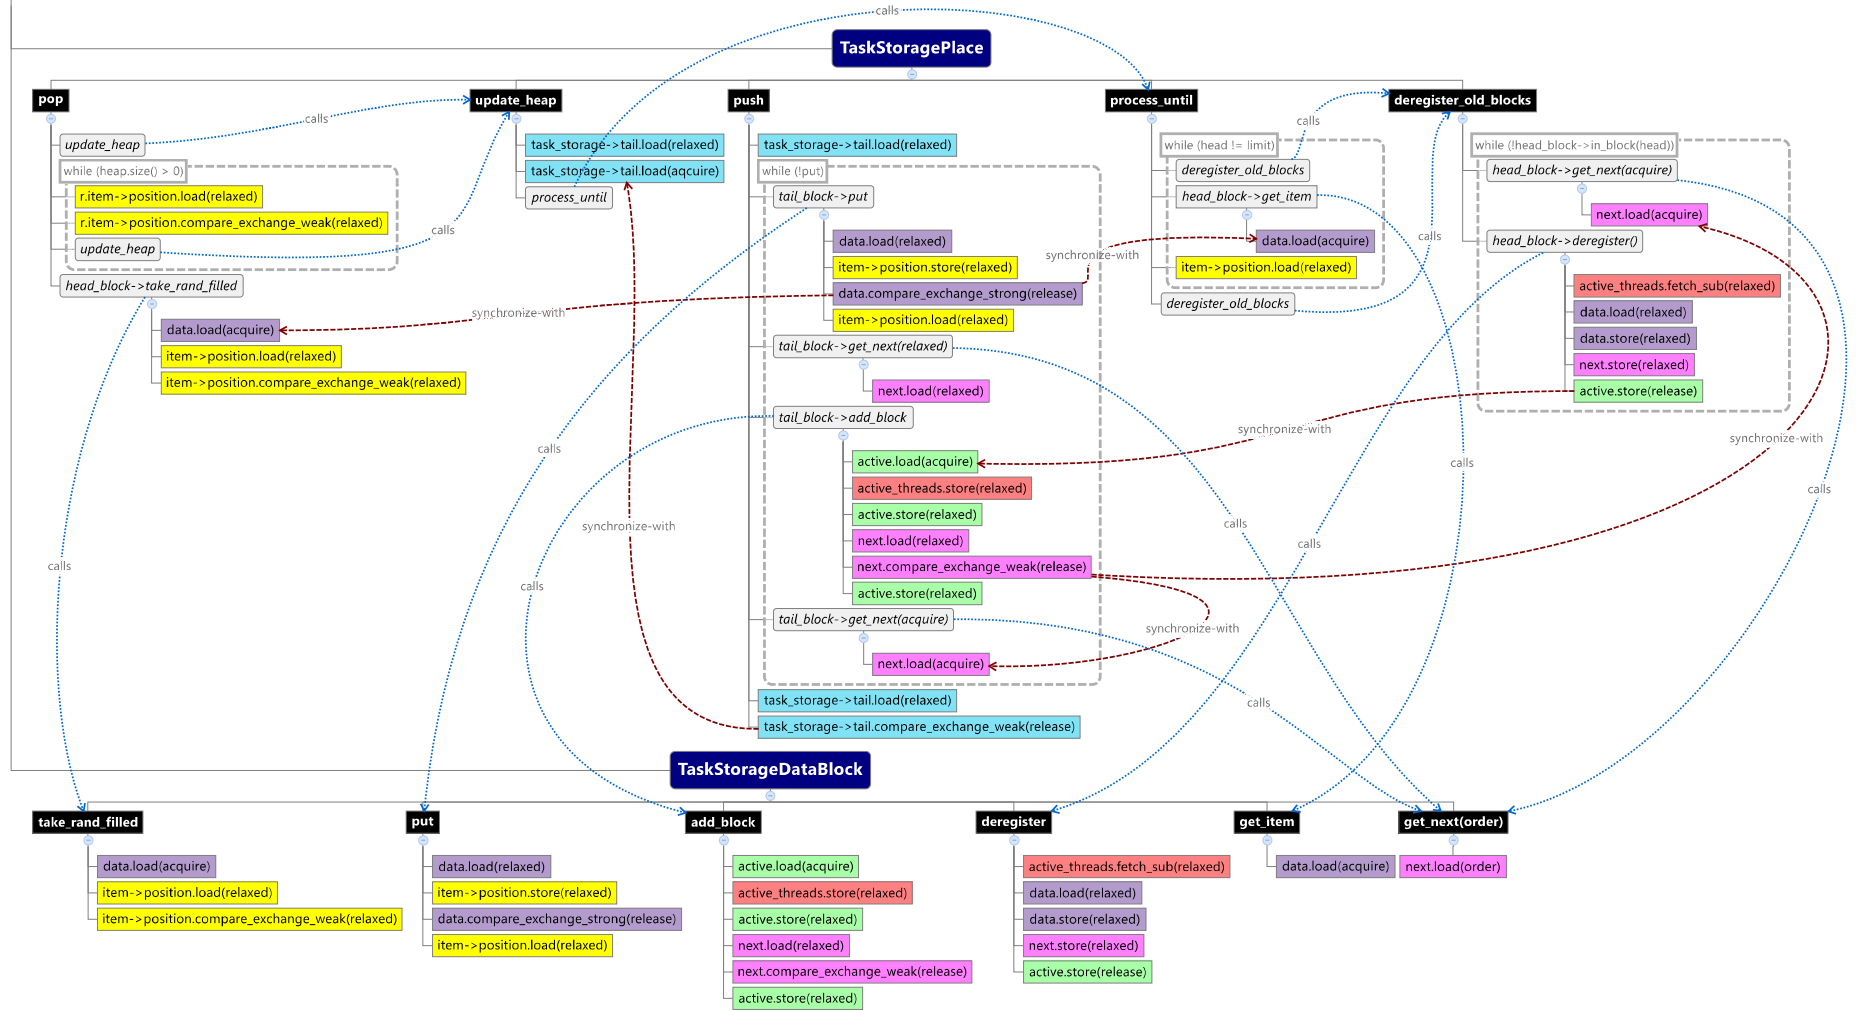


Figure 1

Benchmarks

A set of benchmarks have been run to compare the performance of the original implementation against the new C++11 based implementation. The benchmarks are GP (graph bipartitioning) and SSSP (single source shortest path) from the benchmark suite of the Pheet framework. More information on these and other benchmarks from the Pheet benchmark suite can be found in [10].

They were compiled with gcc 4.9.1 and icc 14.0.1 (Intel C++ Compiler) and run on two machines:

mars

8x Intel® Xeon® E7- 8850 @ 2.00GHz
1TB memory
Linux 3.14-1-amd64 #1 SMP Debian 3.14.4-1 x86_64 GNU/Linux

saturn

4x AMD Opteron™ 6168 @ 800Mhz
128GB memory
Linux 3.2.0-4-amd64 #1 SMP Debian 3.2.46-1 x86_64 GNU/Linux

The benchmarks were run several times with different seeds. The result graphs show the average total runtime.

Graph Bipartitioning

Graph bipartitioning is a well-known, NP-hard problem [7]. Let $G = (V, E)$ be an undirected, weighted graph, $m = |E|$, $n = |V|$. The graph bipartitioning problem is that of partitioning the graph G into two sets with given sizes while minimizing the sum of the weights of the edges that have to be cut.

The Pheet benchmark implements a branch-and-bound approach which is generally well suited to parallelization. It fixes a single node at each step by assigning it to one of the sets for each branch. For bounding (elimination) of sub-problems a simple, easily computable lower bound is used. A more detailed description of this benchmark and its implementation can be found in [10, pp. 146].

The benchmark uses the following parameters:

size the number of nodes in the graph.

p controls the density of the graph. The graph's edges are generated randomly using a symmetric adjacency matrix. For each value of the (half) adjacency matrix an edge is added with probability p .

max_w the maximum weight of an edge. A uniform random distribution between 1 and max_w is used to randomly define the weight of an edge.

The test instance used the following parameter values:

size 35

p 0.9

max_w 1000

It was run 100 times with different seeds.

As can be seen from the results in Figure 2 there is no significant difference in the performance of the two implementations. This is the expected result. The runtime of

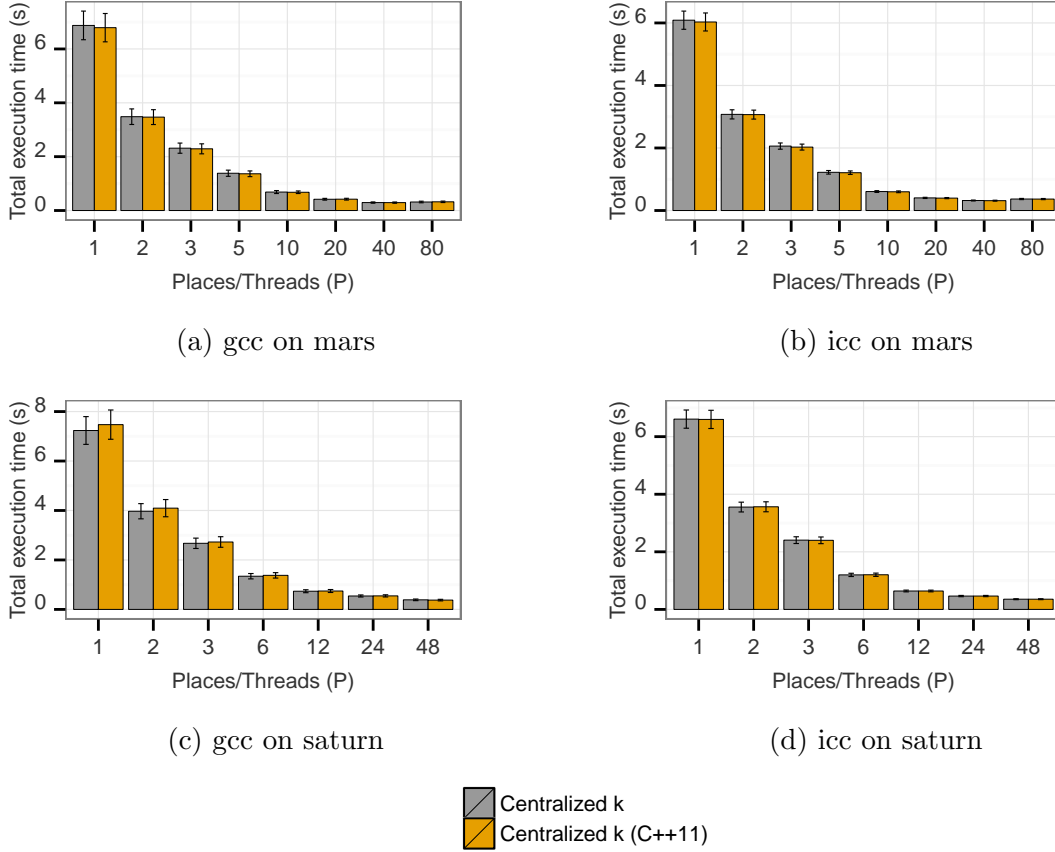


Figure 2: Graph Bipartitioning benchmark results


the individual tasks is relatively high compared to the runtime overhead of the scheduler. Since I only changed some implementation details of the scheduler, but not the scheduling itself these changes can only impact a small amount of the benchmark's total runtime.

Single Source Shortest Path


The shortest path problem is a fundamental and well-studied combinatorial optimization problem with many practical and theoretical applications [9, 2].

Let $G = (V, E)$ be an undirected, weighted graph, $m = |E|$, $n = |V|$, let s be a distinguished vertex of the graph.


The single source shortest path problem (SSSP) is that of computing, for each vertex v reachable from s , a minimum-weight path from s to v , where the weight of a path is the sum of the weights of its edges. The Pheet benchmark implementation is based on a simple parallelization of Dijkstra’s algorithm. A more detailed description of this benchmark and its implementation can be found in [10, pp. 164].



(a) gcc on mars



(b) icc on mars



(c) icc on saturn

Figure 3: Single Source Shortest Path benchmark results for dense graph

The parameters for this benchmark are:

size the number of nodes in the graph.

p controls the density of the graph. The graph’s edges are generated randomly using a symmetric adjacency matrix. For each value of the (half) adjacency matrix an edge is added with probability p .

max_w the maximum weight of an edge. A uniform random distribution between 1 and max_w is used to randomly define the weight of an edge.

k the “relaxation factor“. This is the number of most recent tasks each thread is allowed to miss.

The first instance was run 20 times with different seeds using the following parameters (resulting in a dense graph):

p 0.5
k 512
max_w 100000000
size 10000

The results⁵ in Figure 3 show no significant difference in performance – just as in the graph bipartitioning benchmark.

Because this is a dense graph the generated tasks have high runtime, and therefore again dominate the total runtime. In order to reduce the runtime of the tasks another instance with a sparse graph has been run using the following parameters:

p 0.001
k 1024
max_w 100000000
size 100000

This instance was run 30 times with different seeds. The results are shown in Figure 4.

On Saturn the new C++11 based implementation is significant faster than original implementation, with a p -value of 2.2×10^{-16} according to a paired wilcoxon signed rank test based on the 210 runs ($30 \text{ seeds} \times 7 \text{ thread configurations}$) per implementation. The values for mean runtime and standard deviation for the Saturn benchmark can be found in Table 1.

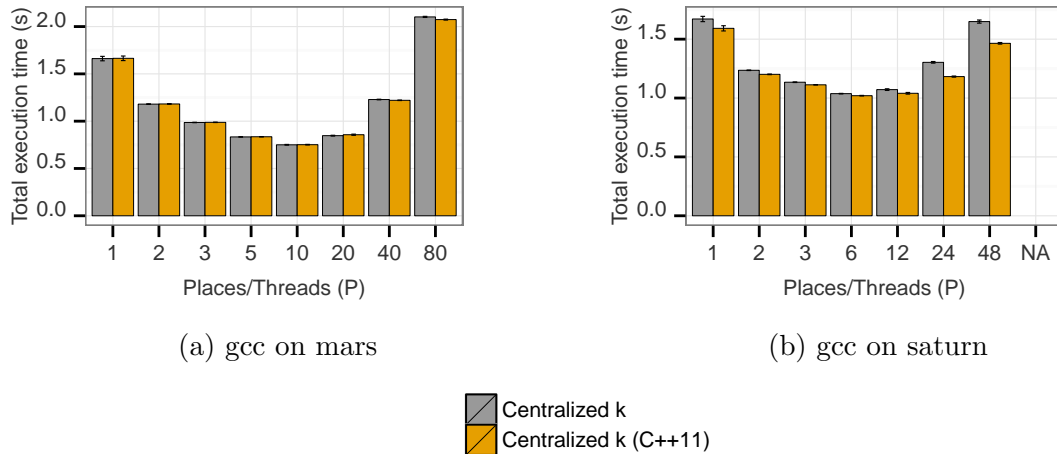


Figure 4: Single Source Shortest Path benchmark results for sparse graph

⁵There are not results for gcc on saturn because the machine was needed by someone else so the benchmark could not be run.

Table 1: Mean runtime (s) and standard deviation for sparse graph on saturn.

Threads	Centralized k		Centralized k (C++11)	
	mean	sd	mean	sd
1	1.6713	0.0601	1.5920	0.0597
2	1.2357	0.0055	1.2013	0.0069
3	1.1343	0.0069	1.1124	0.0048
6	1.0373	0.0060	1.0199	0.0065
12	1.0714	0.0189	1.0403	0.0185
24	1.3032	0.0182	1.1822	0.0125
48	1.6493	0.0358	1.4647	0.0180

The reason for this performance difference most likely is that in the C++11 implementation it is possible to omit an explicit sequential consistent memory fence in `CentralKStrategyTaskStorageDataBlock::deregister`.

Listing 7: Original code of `CentralKStrategyTaskStorageDataBlock::deregister`

```

1 void deregister()
2 {
3     size_t old = SIZE_T_FETCH_AND_SUB(&active_threads, 1);
4     if (old == 1)
5     {
6         // cleanup data items
7         // ...
8
9         next = nullptr;
10        MEMORY_FENCE(); // <—
11        active = false;
12    }
13 }
```

In the original code (shown in Listing 7) this explicit memory fence in line 8 was required to ensure that neither the compiler nor the CPU can reorder the two stores. For the new implementation (shown in Listing 2) these stores have been replaced by two atomic stores – the second one using `memory_order_release`. According to the C++11 standard [6, p. 14] the compiler is not allowed to reorder atomic instructions. It is therefore safe to omit the explicit memory fence.

Looking at the results in Figure 4 it is striking to note that this difference in performance occurs only on Saturn but not on Mars. In the last few years Intel has invested a lot of effort to improve the performance of sequentially consistent operations. The results of the benchmark show that these efforts have really paid off.

Even though it is most likely that the omitting of this memory fence is the cause of the improved scalability I do not have any profiling results that actually prove this assumption.

Conclusion

The Pheet CentralK task storage is a good example to show how different concepts of the new C++11 memory model work and why it is great to have a memory model with clearly defined semantics that allows relaxation for every single operation. The whole data structure is implemented using only acquire/release and relaxed semantics, there is no need for any sequential consistent operations. On some architectures the new implementation even performs better than the original. Making variables atomic and using explicit load/store calls with specific memory order usually produces the same assembler code as the non-atomic operations – at least on x86/x64 architecture. However, it is absolutely necessary to use atomics for variables that are accessed concurrently by multiple threads to avoid data-races. The point of using the new C++11 memory model is not to improve performance but to make the code both - correct and portable. It is simply the only way to write correct (data-race free) and portable multi-threaded C++ code. The resulting C++ code is not only more expressive (atomic variables and their operations can be easily spotted – given that the memory model is always defined explicitly), but also allows formal verification of the correctness – and in some special cases it even results in better performance.

References

- [1] Y. Afek, G. Korland, and E. Yanovsky. Quasi-linearizability: Relaxed consistency for improved concurrency. In *Proceedings of the 14th International Conference on Principles of Distributed Systems*, OPODIS'10, pages 395–410, Berlin, Heidelberg, 2010. Springer-Verlag.
- [2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
- [3] A. Alexandrescu, H.-J. Boehm, K. Henney, D. Lea, and B. Pugh. Memory model for multithreaded C++. Document WG21/N1680=J16/04-0120; available at <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2004/n1680.pdf>, September 2004.
- [4] H.-J. Boehm. Threads cannot be implemented as a library. *SIGPLAN Not.*, 40(6):261–268, June 2005.
- [5] H.-J. Boehm and S. V. Adve. Foundations of the C++ Concurrency Memory Model. *SIGPLAN Not.*, 43(6):68–78, June 2008.
- [6] C++ Standards Committee, Stefanus Du Toit. Working Draft, Standard for Programming Language C++. C++ standards committee paper, <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2012/n3337.pdf>, January 2012.
- [7] M. R. Garey and D. S. Johnson. *Computers and Intractability: a guide to NP-completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.

- [8] S. Meyers and A. Alexandrescu. C++ and The Perils of Double-Checked Locking. *Doctor Dobbs Journal*, Jul 2004.
- [9] R. E. Tarjan. *Data Structures and Network Algorithms*. Society of Industrial and Applied Mathematics (SIAM), 1983.
- [10] M. Wimmer. *Variations on Task Scheduling for Shared Memory Systems*. PhD thesis, Vienna University of Technology, 2014.
- [11] M. Wimmer, D. Cederman, J. L. Träff, and P. Tsigas. Work-stealing with configurable scheduling strategies. In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '13, pages 315–316, New York, NY, USA, 2013. ACM.
- [12] M. Wimmer, M. Poter, and J. L. Traff. The pheet task-scheduling framework on the intel® xeon phi® coprocessor and other multicore architectures. In *Proceedings of the 2013 IEEE 27th International Symposium on Parallel and Distributed Processing Workshops and PhD Forum*, IPDPSW '13, pages 1587–1596, Washington, DC, USA, 2013. IEEE Computer Society.
- [13] M. Wimmer, F. Versaci, J. L. Träff, D. Cederman, and P. Tsigas. Data structures for task-based priority scheduling. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '14, pages 379–380, New York, NY, USA, 2014. ACM.